# Store-Carry-and-Forward Research

William D. Ivancic[+], Wesley Eddy[*], Dennis C. Iannicca[+], Joseph Ishac[+], Alan G. Hylton[+]

## BACKGROUND:

The Store-Carry-and-Forward Research Project was funded by NASA's Office of Chief Technology (OCT) under NASA Glenn Research Center's (GRC) Center Innovation Funds (CIF). This was a 6 month effort. The goals were to research and develop store-carry-and-forward (SCF) technologies for NASA's future communications networks. NASA has identified the need for store-carry-and-forward technologies as critical for its future networks and for interoperability with international partners [1, 2]. The United States Department of Defense is also interested in SCF technologies as evidenced by DARPA funding in Disconnected and Disruption Tolerant Networks [3]. Such technologies are extremely useful for extreme edge environments typical of military operations.

Initial store-and-forward networking studies were performed by the InterPlaNetary Internet Project [4] consisting of a small closed group of avionics software experts at JPL and a few others. The objective of the Interplanetary Internet project was to define the architecture and protocols necessary to permit interoperation of the Internet resident on Earth with other remotely located Internets resident on other planets or spacecraft in transit. In the end, the result was basically a revamping of the monolithic Consultative Committee for Space Data Systems (CCSDS) file delivery protocol (CFDP) into a more layered approach whereby a network shim separated the file delivery application from the transport layer.

Following this effort, a research group was formed in the Internet Research Task Force (IRTF) to investigate this concept. Most of this work was being sponsored by DARPA funding and the European Commission 7th Framework Program for Networking for Communications Challenged Communities (N4C) project [5]. As such the interest moved toward opportunistic networks, military applications, and potential use of multi-level security and content storage to combine disconnected networking with content-based networking. The end result has been the development of an experimental protocol with numerous extensions that have enabled the research community to experiment with a variety of concepts. To date much has been learned using this experimental specification [6,7] Unfortunately, by adopting the CCSDS baseline for disruption tolerant networking, the protocol suite lacks some basic items necessary for a robust protocol. Furthermore, numerous assumptions the have resulted in an inability to secure or scale the protocol. The following are some of the key flaws that need to be corrected in the protocol:

- The current protocol assumes the entire network is time synchronized. This is nearly impossible in a disconnected network and renders the protocol useless for the DOD. In fact, the current Raytheon implementation developed under DARPA funds deviates from the IRTF specification to enable operation in non-synchronized networks.
- There is no checksum in the header. As such, it is impossible to accurately validate proper delivery and storage. The current solution is to use the security protocol specification to provide checksums. However, this means that systems that do not require sophisticated security (e.g. sensorwebs) have to implement it.
- There is no hop-count in the header; data only expires on information lifetime. This has been shown to be a problem already in experimental deployments with bundle storms wasting all available bandwidth. Hop counts are used to expire information due to routing loops.
- The current authentication solution requires an entire "bundle" of information to be received. Since a "bundle" can be any size from a few bytes to Gigabytes, this can result in very easy denial of service attaches as well as resource depletion.
- The security solution proposed makes it impossible to perform reactive fragmentation of large bundles of information. Such reactive fragmentation has proven to be invaluable for distribution of large payloads, which are typical of science data [8].

+ NASA Glenn Research Center
∗ MTI Systems

## INNOVATIONS

The first major innovation we propose is to separate the metadata bundle header from the bundle payload. Such a technique is not typical of normal packet delivery protocols. Today's Internet Protocol (IP) packets (consisting of header and payload) is not overly large (i.e. 1500 bytes for typical packet sent over Ethernet links) and can easily be buffered for forwarding without stressing the router resources – in particular the storage. However, with a store-carry-and-forward (SCF) network, payloads can be extremely large – Megabytes or Gigabytes or more. Due to the large payload size, the combined header and payload is referred to as a "bundle" rather than a packet. Since a SCF router has to look at each bundle before forwarding it and has to do this at the time of, or just prior to, connecting to another SCF router (remember, we are normally disconnected), the need to keep the router (forwarding) tables small is essential. Keeping the metadata header separate from the bundle is a game changer and provides the following advantages:

- The header can be protected separately and, if required, differently than the payload.
- One can authenticate the source of data prior to committing resources. This is extremely important for SCF networks – particularly with regard to thwarting denial-of-service (DOS) attacks in terrestrial use, or, in the case of space systems, only providing resources to those whom you have peering arrangements with.
- Furthermore, one can implement a secure reactive fragmentation process.

The second innovation we propose is to develop an aggregation scheme. Aggregation enables routing tables to be reduced and enables one to determine up front the amount of resources required to accommodate data storage and transfer in a SCF network. By being able to reactively fragment bundles, aggregation becomes practical.

*Justification – Impact and Value:*

Some DTN protocol implementations have large files being transferred using the CCSDS File Delivery Protocol over DTN whereby the file is broken into 64 Kbyte chunks before passing to the DTN agent. This is done for a number of reasons: some do to the protocol design, and some as engineering decisions. Neither is scalable. By limiting bundles to 64 Kbytes, one can easily implement some of the security protocols such as Bundle Authentication. Also, one also does not have to be concerned with pre-allocating large storage for individual bundles. In addition, one does not have to check the bundles size when considering routing algorithms assuming one can transfer at least a single 64 Kbyte bundle in any given contact time. These trades may be fine for command and control, but do not scale and do not apply for science data. Rather, such design decisions de-aggregate science data. For example, if one has a 640 Mbyte payload; the scheme currently in use will de-aggregate this into 10,000 bundles causing route table explosion. That is like ordering a book from Amazon and getting each page sent in a separate envelope. What is really desired is to aggregate whereby if one orders 5 books from Amazon, they are place in a box and then that box is placed on a truck and transported. Each aggregation point reduces the amount of processing and improves the transport efficiency. Such aggregation makes the SCF network scalable. Also, buy aggregating, one knows up front how many resources are being requested by the data source (peering agreement) and can accept/reject based upon authenticated information in the metadata header. This allows one to actually implement and manage true peering agreements.

Recently, DTN implementations have finally realized that the recommendation of splitting large files into 64 Kbyte chunks (bundles) and have been moving away from that recommendation. However, due to the protocol design, authenticating or securing large bundles has opened the system to Denial-of-Service (DOS) attaches as the entire bundle must be received and processed before one can determine if the bundle is going to be accepted. Thus, someone sending bogus multi-megabyte or multi-gigabyte files can easily bring down a DTN agent.

## APPROACH:

Our roadmap for this research and development had three major processes:

1. Develop a problem statement and requirements document and publish
2. Develop the Store-Carry-and-Forward protocol specification and publish
3. Implement and test the specification and publish the results

Our approach to addressing SCF research and development was to take a fresh look at the problem. Our primary goal was to define the problem such that it was generic and common to many situations, not just, for example, deep space communications. This was followed by the generation of a requirements document and test requirements document. These items were critical because, with only a 6-month effort, the documents would allow others to pick up where we left off. Once we had the problem and requirements well documented, our next step was to the naming and addressing aspects of SCF. Naming and addressing have, to date, been implemented poorly resulting in architectures that are insecure and do not handle mobility and multi-homing well. Once naming and addressing was completed, we would move on to protocol development and finally implementation and testing.

## ACCOMPLISHMENTS

We successfully completed development of the problem statement and two requirements documents. These have been documented as Internet Engineering Task Force (IETF) Internet Drafts (IDs). An explanation of the Internet standards process is provided in RFC 2026 [9]. An ID is a draft version of the document that may become an informational document or a protocol specification. A document that describes a "best current practices" document, a requirements document or problem statements document are informational documents. A document that describes a protocol is a protocol specification. In addition, we studied and documented naming and addressing schemes that could be used for SCF and general systems. This work has been submitted for conference publication. Notice of acceptance or rejection is scheduled for December 2012.

Our problem statement document is entitled "Store, Carry and Forward Problem Statement [10]." This document provides a problem statement for non-realtime communication between systems that are generally disconnected, requiring multiple network hops between the source and destination that may never be fully connected end-to-end at any given time. This document describes a number of use cases that motivate having a standard method to communicate between such systems, as multi-organization and multi-vendor support and interoperability is highly desirable. These include dismounted soldiers, sensorwebs, medical devices, animal tracking, low-earth-orbiting satellites and data mule scenarios. To avoid confusion in terminology when trying to focus on the problem and requirements without bias towards particular technical solutions, at this time, we refer to the protocol instances that would support such communications as Store, Carry, and Forwarding (SCF) agents, and refer to their activity as SCF networking. The concepts involved in SCF networking are not entirely new and several facets of the problem have been solved in multiple incompatible ways in existing or historic systems. This document describes the core SCF problem and gives an assessment of the ability to use existing technologies as solutions.

The second document we generated was SCF requirements and expectations document, "Store, Carry and Forward Requirements and Expectations [11]." This document describes the requirements for a Store, Carry and Forward (SCF) protocol, and the expectations placed upon the SCF agents and SCF applications. The document includes sections on: Design Considerations, Protocol Requirements, SCF Agent Requirements and Expectations, Application Requirements and Expectations, and Security Considerations. It was important to include SCF operation expectations as this has been notably neglected in the past and results in systems that do not interoperate well, not because of the protocol, but because of the expectation of what various entities will handle information. For example, if a forwarding agent does not have a current route to a destination, what should it do: drop the information and say nothing, drop the information and inform the previous hop, or hold and wait for some time to see if a route gets established? None of these are wrong, so long as everyone understands what is expected.

The third document is a currently just a framework for testing requirements, "Store, Carry and Forward Testing Requirements [12]."This document provides guidelines and requirements for testing Store, Carry and Forward (SCF) systems and protocols. The primary motivation for developing this document is to establish thorough, repeatable tests that will fully exercise a SCF system. Past experience has shown that testing of SCF systems to often be inadequate. For example, tests have been performed on SCF systems in fully connected, high bandwidth networks where only forwarding would be exercised or the traffic would be so minimal as to never tax the storage or queuing. Such tests are valid as a starting point, but insufficient to determine that a protocol or implementation will work properly in a reasonably scaled deployment. A secondary motivation is to improve implementations by providing a known test environment. Knowing some possible ways that the protocol and system will be evaluated may help establish how the code is developed as well as identifying hooks for monitoring particular processes.

The naming and addressing paper is entitled, "Secure Naming and Addressing Operations for Store, Carry and Forward Networks." This paper describes concepts for secure naming and addressing directed at Store, Carry and Forward (SCF) networks, networks where disconnection is the norm. The paper provides a brief overview of store, carry and forward networks followed by an in depth discussion of how to securely: create a namespace; allocate names within the namespace; query for names known within a local processing system or connected subnetwork; validate ownership of a given name; authenticate data from a given name; and, encrypt data to a given name. Critical issues such as revocation of names, mobility and the ability to use various namespaces to secure operations or for Quality-of-Service are also presented. Although the concepts presented for naming and addressing have been developed for SCF, they are directly applicable to fully connected networks [13].

## TECHNICAL DISCLOSURES

There are no patent disclosures for this work nor will there likely be a patented protocol. If one wishes for a protocol to be used, it must have open standards and be freely available. Publishing is sufficient to keep anyone else from patenting the idea and ensuring the protocol is available, free of charge, to anyone who wishes to deploy it.

The quickest and best way to disclose the work is via publishing as Internet Drafts. This is what we did. Internet drafts are available from the IETF servers to anyone wishing to obtain a copy. Furthermore, the IETF is an open standards body of individuals. It is not a political organization or company-run. Anyone can participate. There is no cost to be a member and most work is done via email lists with all meetings open to anyone who is interested. All meetings are broadcast over the web via MP3 streaming and remote participation is available via streaming XML protocol (XMPP) Instant Messaging.

We also disclosed this work to the various groups of the Department of Defense, and their contractors and university parities. On June 26, 2012, Mr. William Ivancic and Mr. Alan Hylton participated in the Robust, Distributed Networking Transport Workshop at Aberdeen Proving Grounds in Aberdeen, Maryland. The workshop was co-sponsored by the US Army Communications-Electronics Research, Development and Engineering Center (CERDEC) and the Naval Research Laboratories (NRL). Mr. Ivancic presented the current state of his research on Store, Carry and Forward (SCF) systems including: a problem statement, scenarios, operation expectations of SCF agents, and requirements. The overreaching theme of the workshop was efficiently getting information to the war fighter (the edge of the network). Technologies include: Delay Tolerant Networking (DTN), mobile ad hoc networking (manet), intelligent caching, and making client server applications such as Chat and Bit Torrent into distributed applications.

## UNIVERSITIES & STUDENTS ENGAGED

Currently, no students have been directly engaged. However, we have engaged John Day of Boston University to review our work as much of our concepts have been based on some of his initial work related to Inter Process Control concept regarding networks and layering of networks. We have received the feedback and hope to continue working with John and his team at Boston University.

## IMPROVING THE CULTURE OF INNOVATION AT GLENN RESEARCH CENTER

The CIF funds have definitely enabled creative thought and innovation at GRC. There is no question, that without such funding and without the ability of the Deputy Chief Technologist to independently evaluate and select proposals based on their technical content, the work reported here would never have been done. In this respect, the CIF Funds have definitely improved the culture of innovation at GRC.

The CIF funds have also allowed researchers at GRC to provide the Robust, Distributed Networking Transport research community and US Army Communications-Electronics Research, Development and Engineering Center (CERDEC) a new perspective on Store, Carry and Forwarding techniques and protocols. This should further the innovation of these concepts and technologies within the global research community.

Note well: This work has done very little to improve the culture of innovation within NASA's Space Communication and Navigation (SCaN) Program at GRC or HQ. Through NASA Policy Directive NPD 8074.1 [14], SCaN has become the gatekeeper for all space communication research. SCaN's practices have show they consider this authority to extend to anything that SCaN considers may affect SCaN's program even if it is not related to space communications. Hence, programs within the Science Mission Directorate such the Earth Science Technology Office (ESTO) are reluctant to fund communication research even for sensorwebs or airspace systems. Furthermore SCaN has attempted to and continues to (with some success) thwart open discussion and collaboration on communication networking, architecture and protocol development with various research communities. Communication work related to network architecture, and protocol development has become extremely political to the detriment of sound technical engineering let alone research. Such practices, most assuredly, discourage innovation regarding communication networking, architecture and protocols.

## TECHNOLOGY IMPACT

In the Internet, there are two namespaces, IP addresses for routing, and DNS names for higher-level identifiers. The limitation of two namespaces, and the global visibility of those namespaces, is a root cause of many complexities and fragilities within today's Internet architecture, including within: the interdomain routing system, the Domain Name Services (DNS), IP neighbor discovery, and other aspects. This has led to a multitude of security issues related to not being able to verify ownership of particular identifiers or addresses, and not being able to authenticate the bindings between particular identifiers and addresses.

For SCF, we are proposing a system of unlimited namespaces, which can be used to construct either pools of identifiers without mandated structure, or pools of addresses with hierarchical structure. The only difference between addresses and other identifiers is their hierarchical nature. The secure naming system developed under this research task provides a light-weight method for allocating and validating application names and locators (addresses) that could be deployed in a Store, Carry and Forward, normally disconnected networks. The technique can also be applied to fully connected networks. By ensuring that the application names separate from the location names, the system readily handles multi-homing and mobility.

Our system could be an enabling technology for the aeronautics networks vastly simplifying operations and management. For instance, every infrastructure provider can maintain its own namespaces for management of its equipment. Since these are not exposed to the users, most security threats to the infrastructure instantly disappear.

Infrastructure providers that wish to confederate for the purposes of creating a routable address space between them can do so, and those routable addresses still do not expose their management and control planes to one another. Mobile users sharing Namespace Identifier (NSI) certificates for that address space can roam to any provider that's also part of the namespace and obtain addresses without any pre-existing trust relationships. If they need to be globally reachable themselves, they can use their own namespaces created for specific domains and allowing applications from all domains to utilize the same infrastructure yet be completely isolated from one another except for sharing bandwidth. Such techniques also apply to

William I
**Deleted:**
the culture
to how NA
related wor
attitude tow
Communic
has made th
communica
consider ma
the Science
communica
research, bu
Also, SCaN
funding wo
Technology
not care to
discourages
discussion
with variou
become ext
technical en

William I
**Deleted:**
confederate
address spa

securing "Critical Infrastructure Networking". There will be no fear of accidentally leaking routes because the namespaces have been factored out, access to names is secured, and proof of ownership is verified.

## DESIRED NEXT STEPS

The Center Innovation Funding is all that is currently available for this research within NASA at this time. Thus, the 1st goal is to get as much information out for others to build on. Graduate students and their professors are constantly searching for interesting research that may have an impact.

The desired next step is to obtain support within NASA or the US Government to continue this research *at NASA GRC.*

We have established a relationship with the US Army Communications-Electronics Research, Development and Engineering Center (CERDEC). We have some money for travel and to aid CERDEC in investigating how to get information to the edge of the network, the war fighter. Technologies include: Delay Tolerant Networking (DTN), mobile ad hoc networking (manet), intelligent caching, making client server applications such as Chat and Bit Torrent into distributed applications, and SCF. We may be able to leverage our relationship with CERDEC to obtain additional funding. The desire is to be able to perform a detail with CERDEC to help them with some immediate needs as those needs closely mirror NASA's future needs in areas such as sensorwebs and for space networking.

Secure naming has the potential to simplify aeronautical communication. It enables multi-homing, mobility, and sharing of infrastructure thereby reducing overall system costs. Air Traffic Control (ACT), Airline Operations Control (AOC) and Passenger Entertainment and Internet Services (PIES) can all securely share the same infrastructure when maintaining secure operation by having each separate domains applications operate in separate namespaces. Thus, funding for at least the naming and addressing portion of SCF may be available via the aeronautics program.

Secure naming and addressing is applicable to securing "Critical Infrastructure Networking". Thus, there is likely to be funding opportunities in this area, but not from NASA even though NASA has "Critical Infrastructure". One may have to leave NASA to pursue such opportunities. This certainly is an area GRC could investigate to bring new business to the center.

William I
Deleted:
William I
Deleted:
William I
Formatte
Roman,
William I
Deleted:
thus, the sp
William I
Deleted:
William I
Deleted:

---

1 "Integrated Network Architecture Definition Document Volume 1," January 28. 2010

2 Charles D. Edwards (JPL), Michel Denis (European Space Operations Centre), Lena Braatz (Booz Allen Hamilton): "Concept for a Solar System Internetwork," IEEE Aerospace Conference, March 2011

3 BAA04-13 - Disruption Tolerant Networking (DTN), May 2005

4 InterPlaNetary Internet Project (IPN), http://www.ipnsig.org/home.htm, March 2011

5 http://www.n4c.eu/N4Cproject.php, March 2011

6 V. Cerf et al., "Delay-Tolerant Network Architecture," Internet Engineering Task Force (IETF) RFC 4838, informational, April 2007.

7 K. Scott, and S. Burleigh, Bundle Protocol Specification, Internet Engineering Task Force (IETF) RFC 5050, experimental, November 2007.

8 Will Ivancic, Phil Paulsen, Dave Stewart, John Taylor, Scott Lynch, Jay Heberle, James Northam, Chris Jackson and Lloyd Wood, "Large File Transfers from Space using Multiple Ground Terminals and Delay-Tolerant Networking," IEEE Global Communications Conference (Globecom 2010), Miami, Florida, December 2010

9 S. Bradner, "The Internet Standards Process -- Revision 3," RFC 2026, BCP: 9, October 1996

10 W. Ivancic, W. Eddy, A. Hylton, D. Iannicca, J. Ishac, "Store, Carry and Forward Problem Statement," draft-ivancic-scf-problem-statement-00, July 9, 2012 work-in-progress http://tools.ietf.org/id/draft-ivancic-scf-problem-statement-00.html

11 W. Ivancic, W. Eddy, A. Hylton, D. Iannicca, J. Ishac, "Store, Carry and Forward Requirements and Expectations," draft-ivancic-scf-requirements-expectations-00, July 9, 2012 work-in-progress

http://tools.ietf.org/id/draft-ivancic-scf-requirements-expectations-00.html

[12] W. Ivancic, W. Eddy, A. Hylton, D. Iannicca, J. Ishac, "Store, Carry and Forward Testing Requirements," draft-ivancic-scf-testing-requirements -00, July 9, 2012 work-in-progress http://tools.ietf.org/id/draft-ivancic-scf-testing-requirements- 00.html

[13] W. Eddy, W. D. Ivancic, D. C. Iannicca, J. Ishac, A. G. Hylton. "Secure Naming and Addressing Operations for Store, Carry and Forward Networks," submitted to for the 10th USENIX Symposium on Networked Systems Design and Implementation. Lombard, IL, April 3–5, 2013

[14] NASA Policy Directive NPD 8074.1, Effective Date: August 11, 2009, Expiration Date: August 11, 2014